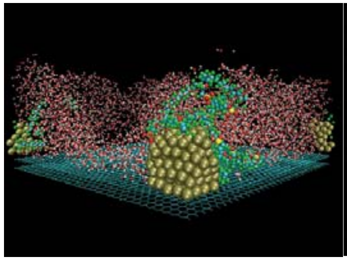 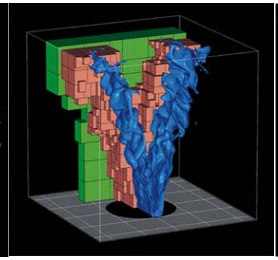 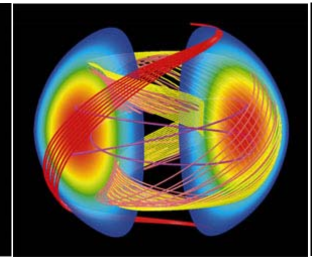 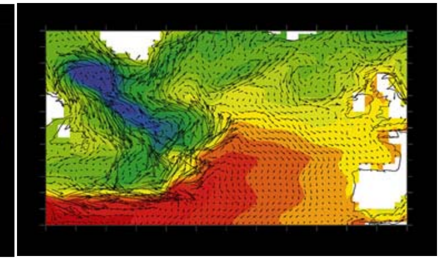 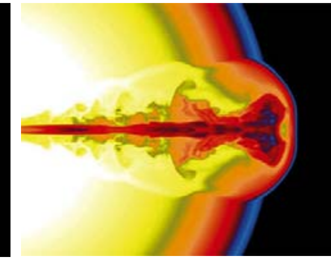

Towards High Sustained Applications Performance on Leadership Computer Systems

Horst Simon
Lawrence Berkeley National Laboratory

SC04 Pittsburgh,

November 10, 2004

BOF on Leadership Computing Systems

Scientific Applications and Underlying Algorithms Drive Architectural Design

- Requirement: 50 Tflop/s sustained in 2007 on applications of national importance
- Process:
  - identify applications
  - identify computational methods used in these applications
  - identify architectural features most important for performance of these computational methods

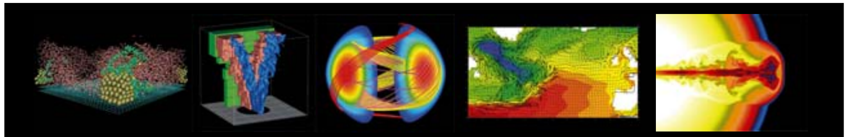

Leadership Computing Applications in DOE

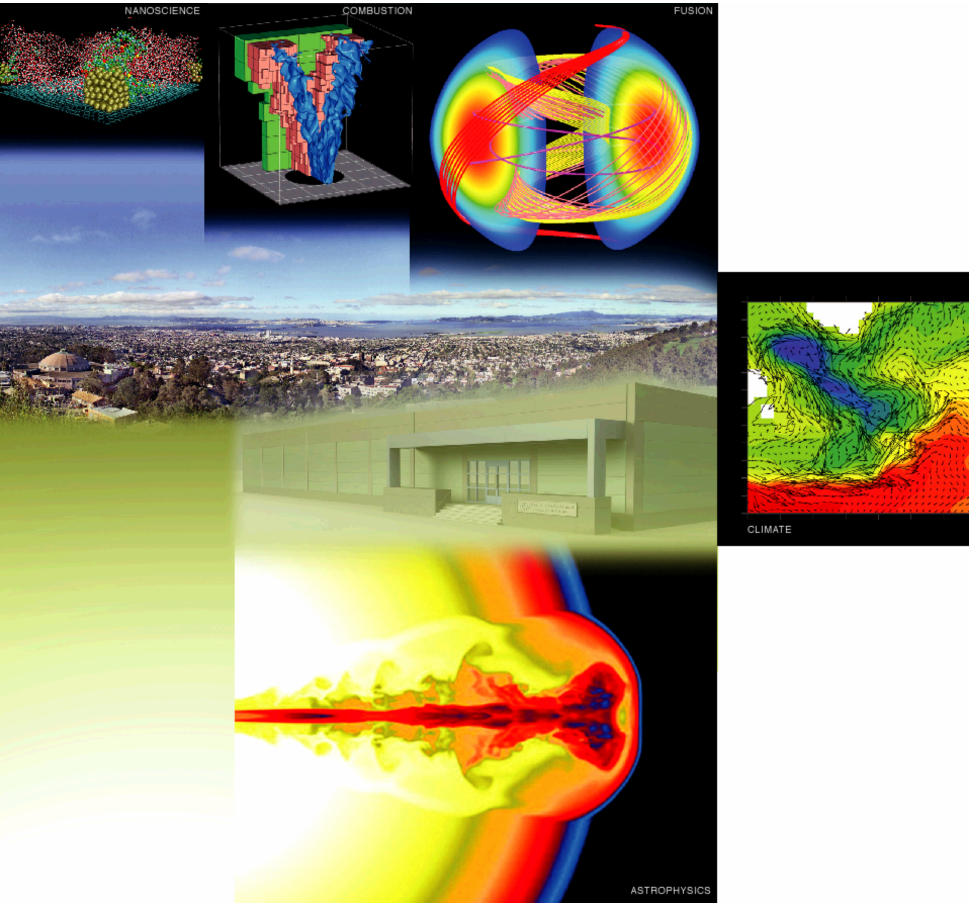

- Nanoscience
- Combustion
- Fusion
- Climate modeling
- Astrophysics

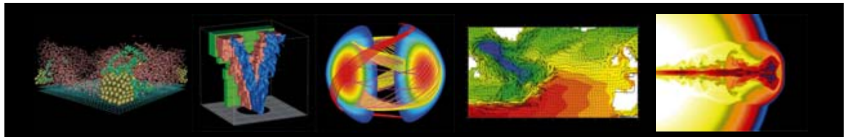

Science Breakthroughs Enabled by Leadership Computing Capability

Science Areas	Goals	Computational Methods	Breakthrough Target (50Tflop/s)
Nanoscience	Simulate the synthesis and predict the properties of multi-component nanosystems	Quantum molecular dynamics Quantum Monte Carlo Iterative eigensolvers Dense linear algebra Parallel 3D FFTs	Simulate nanostructures with hundreds to thousands of atoms as well as transport and optical properties and other parameters
Combustion	Predict combustion processes to provide efficient, clean and sustainable energy	Explicit finite difference Implicit finite difference Zero-dimensional physics Adaptive mesh refinement Lagrangian particle methods	Simulate laboratory scale flames with high fidelity representations of governing physical processes
Fusion	Understand high-energy density plasmas and develop an integrated simulation of a fusion reactor	Multi-physics, multi-scale Particle methods Regular and irregular access Nonlinear solvers Adaptive mesh refinement	Simulate the ITER reactor
Climate	Accurately detect and attribute climate change, predict future climate and engineer mitigation strategies	Finite difference methods FFTs Regular and irregular access Simulation ensembles	Perform a full ocean/ atmosphere climate model with 0.125 degree spacing, with an ensemble of 8-10 runs
Astrophysics	Determine through simulations and analysis of observational data the origin, evolution and fate of the universe, the nature of matter and energy, galaxy and stellar evolutions	Multi-physics, multi-scale Dense linear algebra Parallel 3D FFTs Spherical transforms Particle methods Adaptive mesh refinement	Simulate the explosion of a supernova with a full 3D model

Berkeley Lab

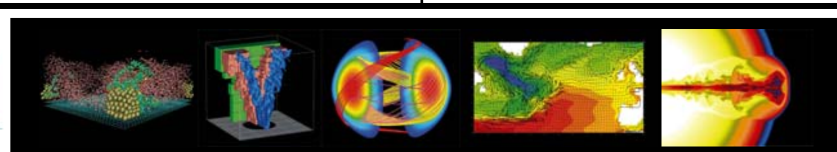

Office of Science
U.S. Department of Energy

Science-Driven Technology

Recent trend:

- Systems designed for applications other than high performance computing (web servers, desktop applications, databases, etc.)

Result:

- Scientists are increasingly frustrated by low sustained performance rates on real-world, high-end problems (typically 5-10% of advertised peak performance)

What is needed:

- A science-driven system architecture that achieves high sustained performance on a broad range of scientific applications.

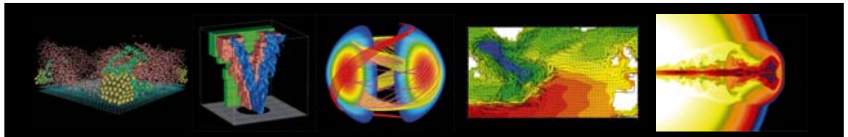

How Science Drives Architecture

State-of-the-art computational science requires increasingly diverse and complex algorithms

Science Areas	Multi-Physics and Multi-Scale	Dense Linear Algebra	FFTs	Particle Methods	AMR	Data Parallelism	Irregular Control Flow
Nanoscience	X	X	X	X		X	X
Combustion	X			X	X	X	X
Fusion	X	X		X	X	X	X
Climate	X		X		X	X	X
Astrophysics	X	X	X	X	X	X	X

Only balanced systems that can perform well on a variety of problems will meet future scientists' needs!

Data-parallel and scalar performance are both important

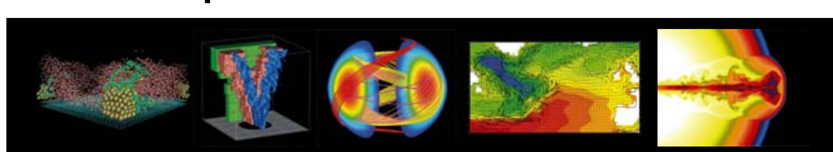

“Evaluation of Leading Superscalar and Vector Architectures for Scientific Computations”

Leonid Oliker, Andrew Canning, Jonathan Carter

LBNL

Stephane Ethier

PPPL

(see paper in SC04 technical program)

Berkeley Lab

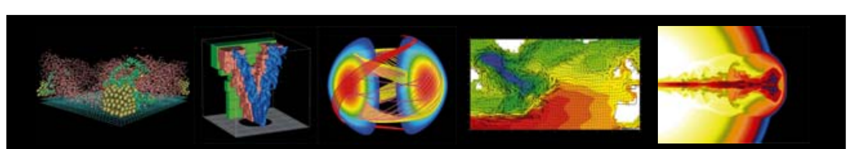

Office of Science
U.S. Department of Energy

PARATEC: Performance

Data Size	P	Power 3		Power4		Altix		ES		X1	
		Gflops/P	%peak	Gflops/P	%peak	Gflops/P	%peak	Gflops/P	%peak	Gflops/P	%peak
432 Atom	32	0.95	63%	2.0	39%	3.7	62%	4.7	60%	3.0	24%
	64	0.85	57%	1.7	33%	3.2	54%	4.7	59%	2.6	20%
	128	0.74	49%	1.5	29%	---	---	4.7	59%	1.9	15%
	256	0.57	38%	1.1	21%	---	---	4.2	52%	---	---
	512	0.41	28%	---	---	---	---	3.4	42%	---	---
686 Atom	128							4.9	62%	3.0	24%
	256							4.6	57%	1.3	10%

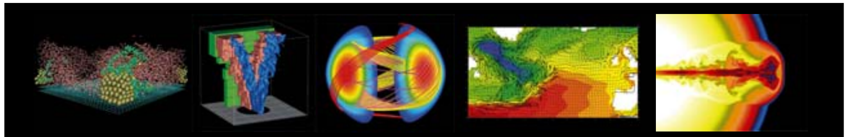

GTC: Performance

Number Particles	P	Power 3		Power4		Altix		ES		X1	
		Gflops/P	%peak	Gflops/P	%peak	Gflops/P	%peak	Gflops/P	%peak	Gflops/P	%peak
10/cell 20M	32	0.13	9%	0.29	5%	0.29	5%	0.96	12%	1.00	8%
	64	0.13	9%	0.32	5%	0.26	4%	0.84	10%	0.80	6%
100/cell 200M	32	0.13	9%	0.29	5%	0.33	6%	1.34	17%	1.50	12%
	64	0.13	9%	0.29	5%	0.31	5%	1.25	16%	1.36	11%
	1024	0.06	4%								

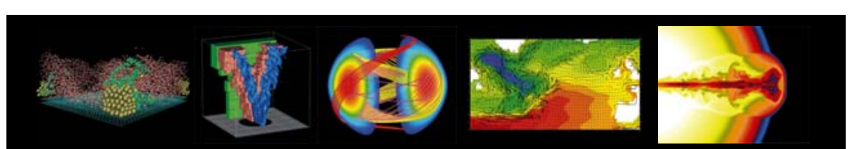

New Science Presents New Architecture Challenges

Leadership computing requires an architecture capable of achieving high performance across a spectrum of key state-of-the-art applications.

- Data parallel algorithms do well on machines with high memory bandwidth (vector or superscalar)
- Irregular control flow requires excellent scalar performance
- Spectral and other methods require high bisection bandwidth

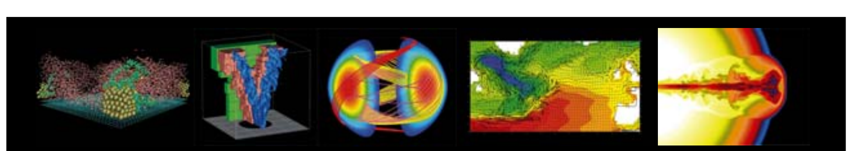

Scalar Performance Increasingly Important

- Cannot use dense methods for largest systems because of N^3 algorithm scaling. Need to use sparse and adaptive methods with irregular control flow
- Complex microphysics results in complex inner loops

"It would be a major step backward to acquire a new platform that could reach the 100 Tflop level for only a few applications that had 'clean' microphysics. Increasingly realistic models usually mean increasingly complex microphysics. Complex microphysics is not amenable to [simple vector operations]."

– Doug Swesty, SUNY Stony Brook

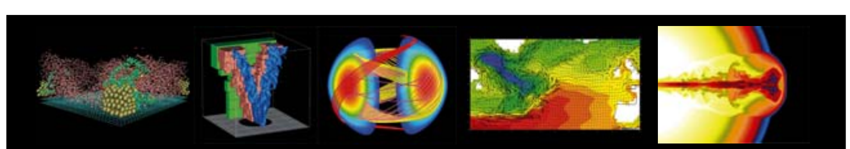

DARPA HPCS will characterize applications

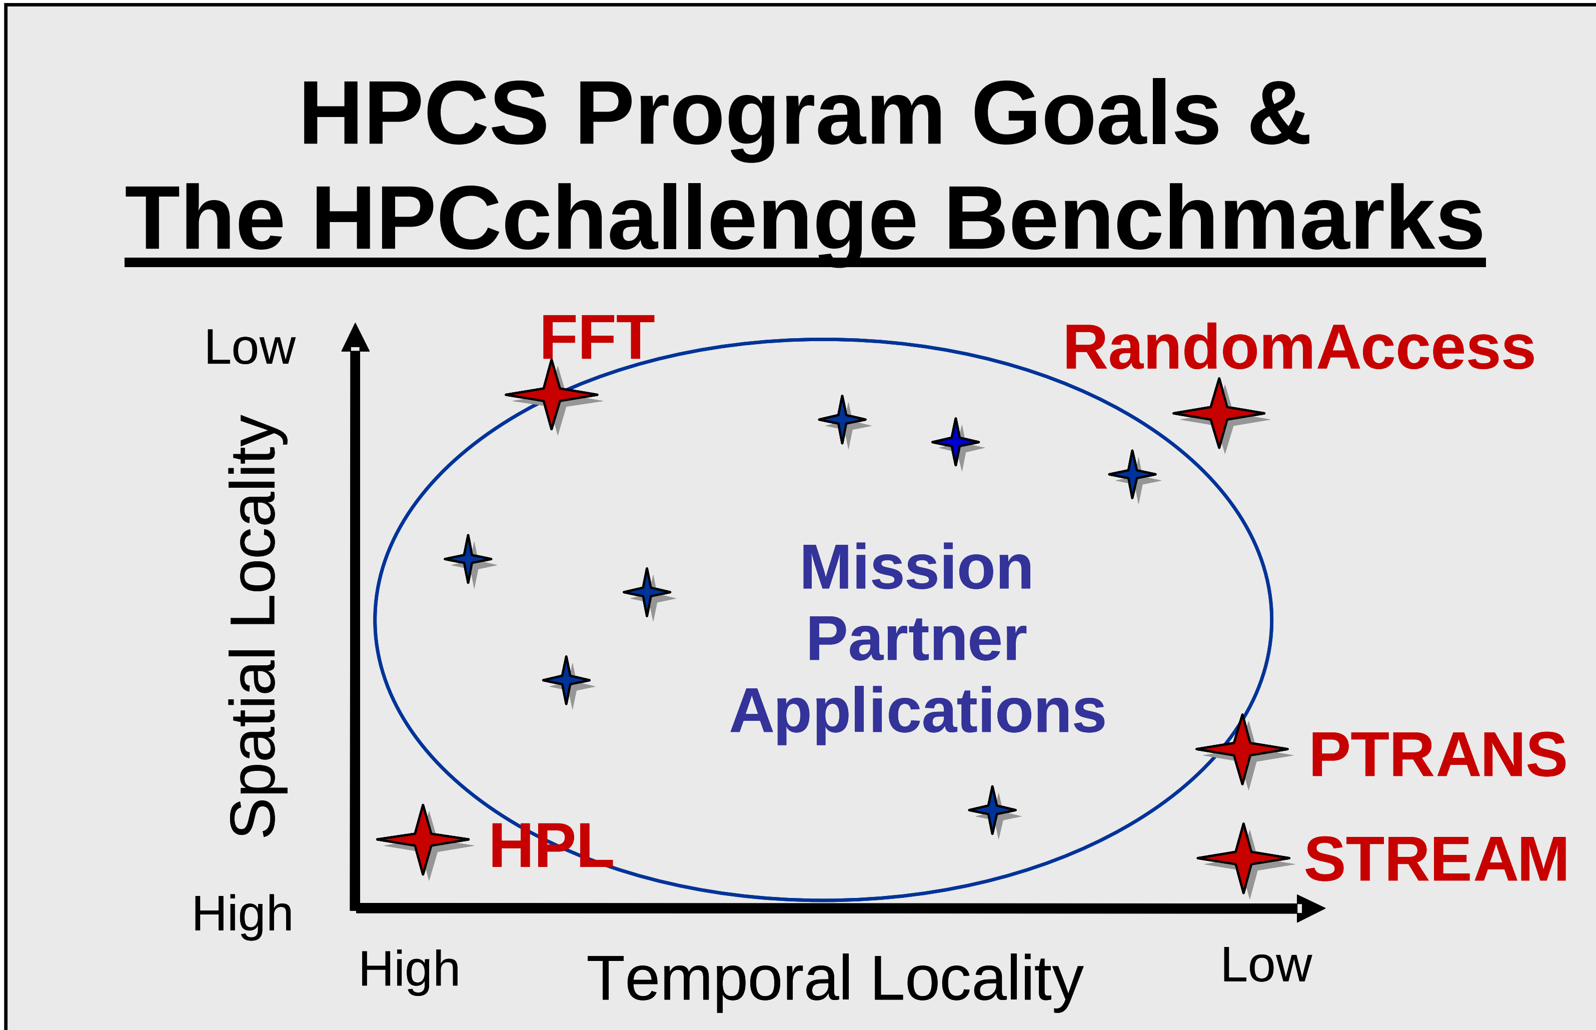

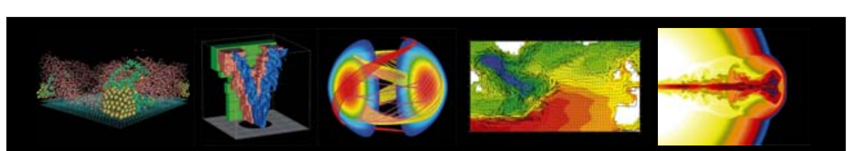

APEX-Map: A Synthetic Benchmark to Explore the Space of Application Performances

Erich Strohmaier, Hongzhang Shan
Future Technology Group, LBNL
EStrohmaier@lbl.gov

Co-sponsored by DOE/SC and NSA

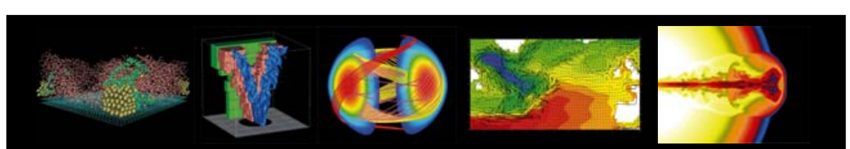

Apex-MAP characterizes architectures through a synthetic benchmark

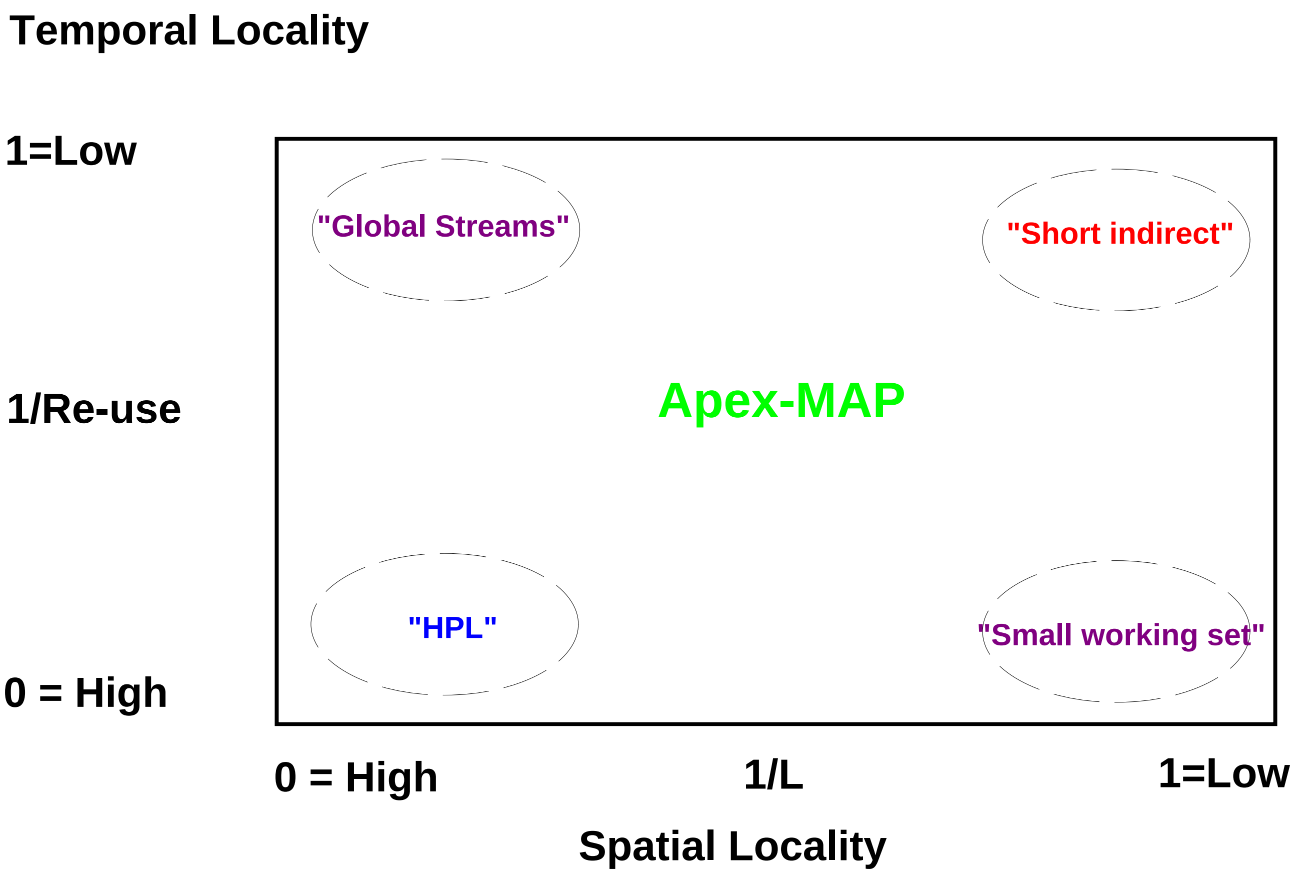

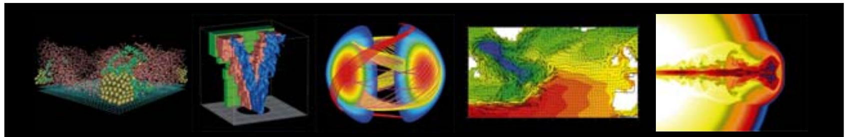

Apex-Map Sequential

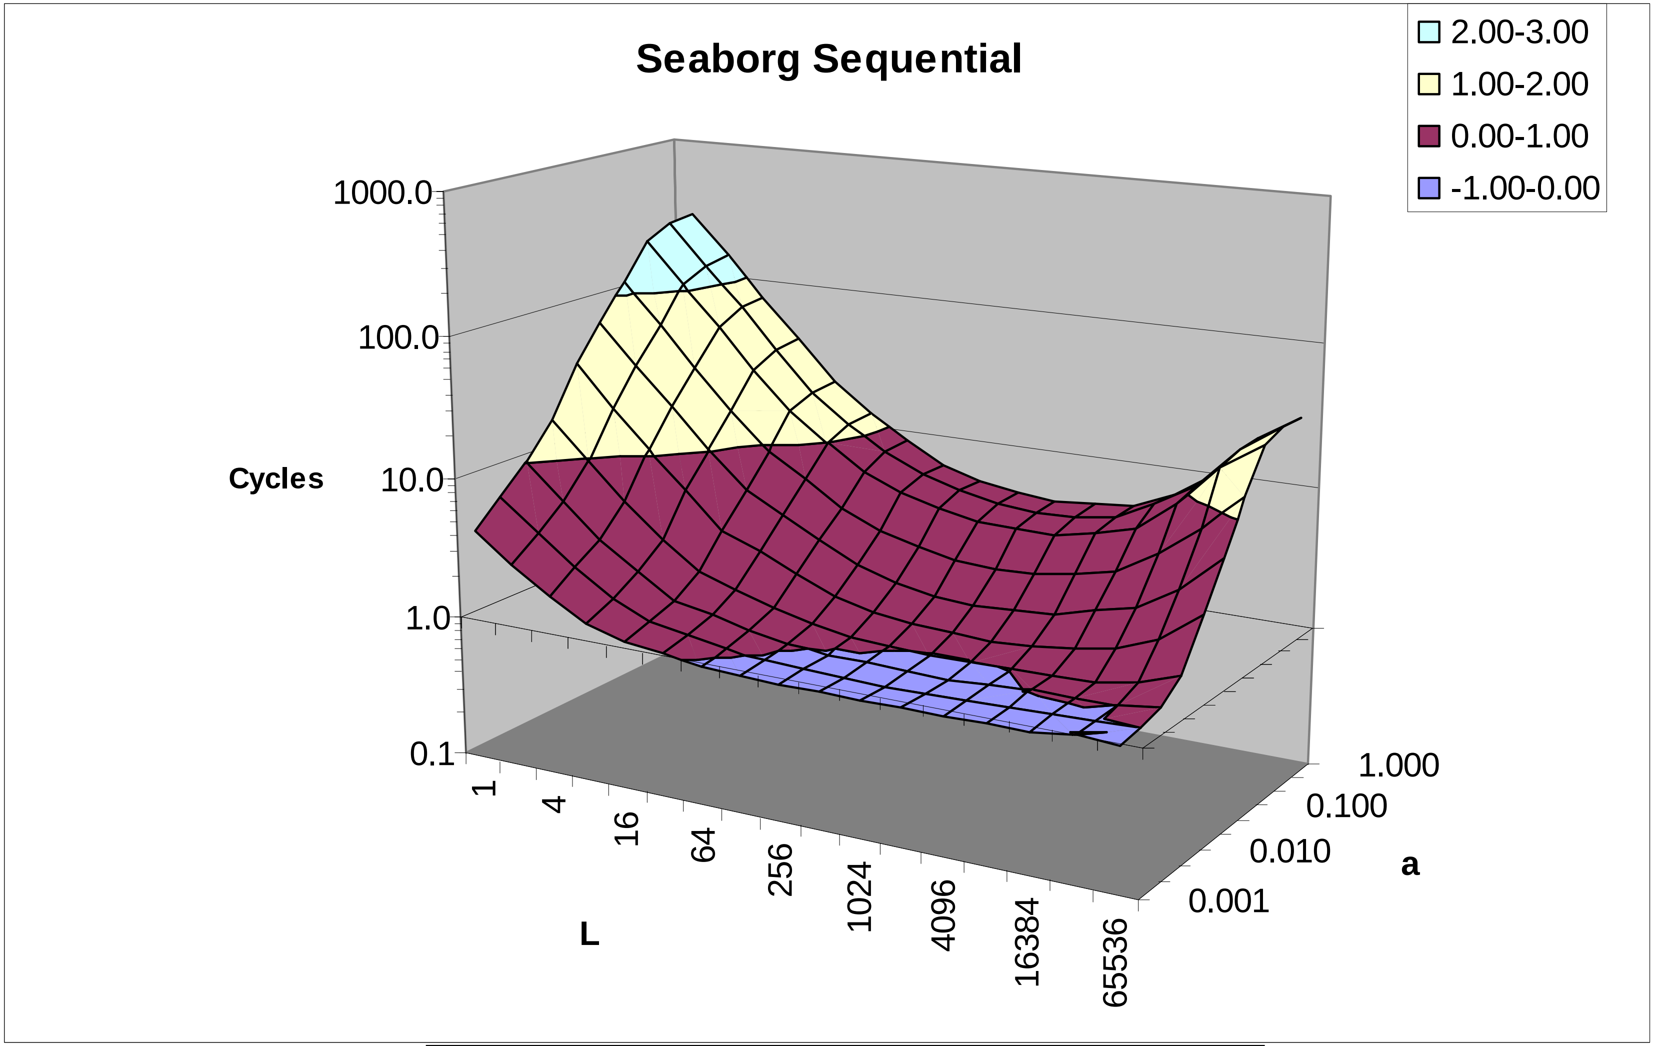

Apex-Map Sequential

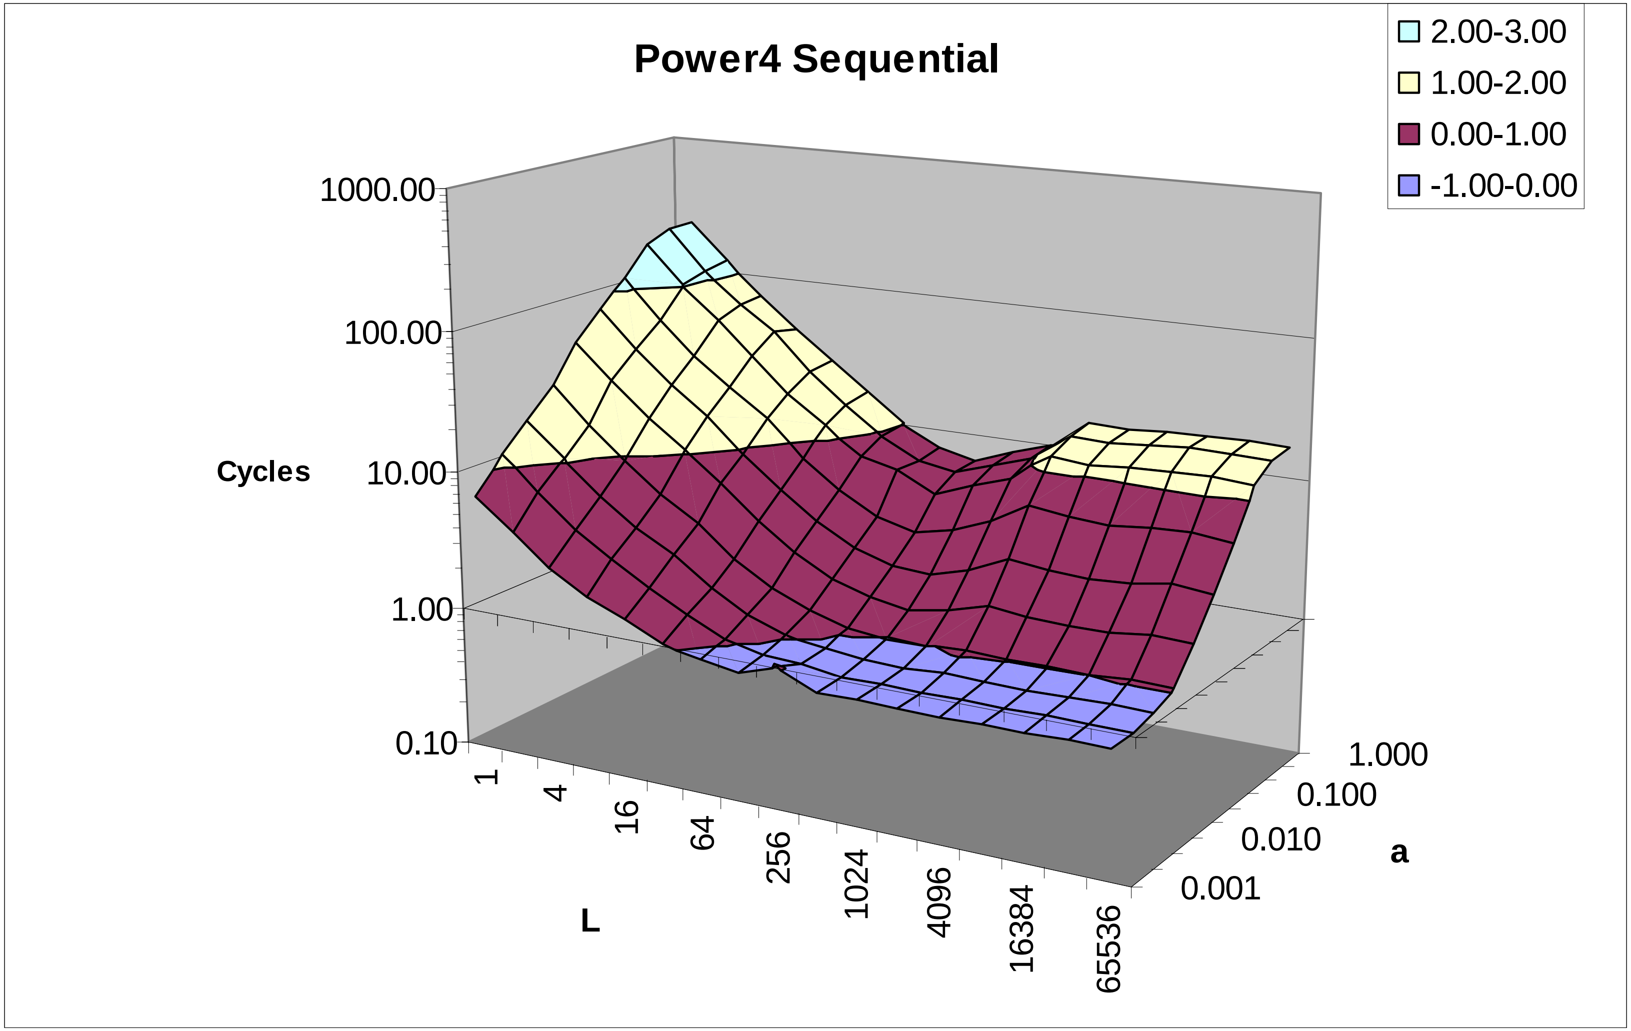

Apex-Map Sequential

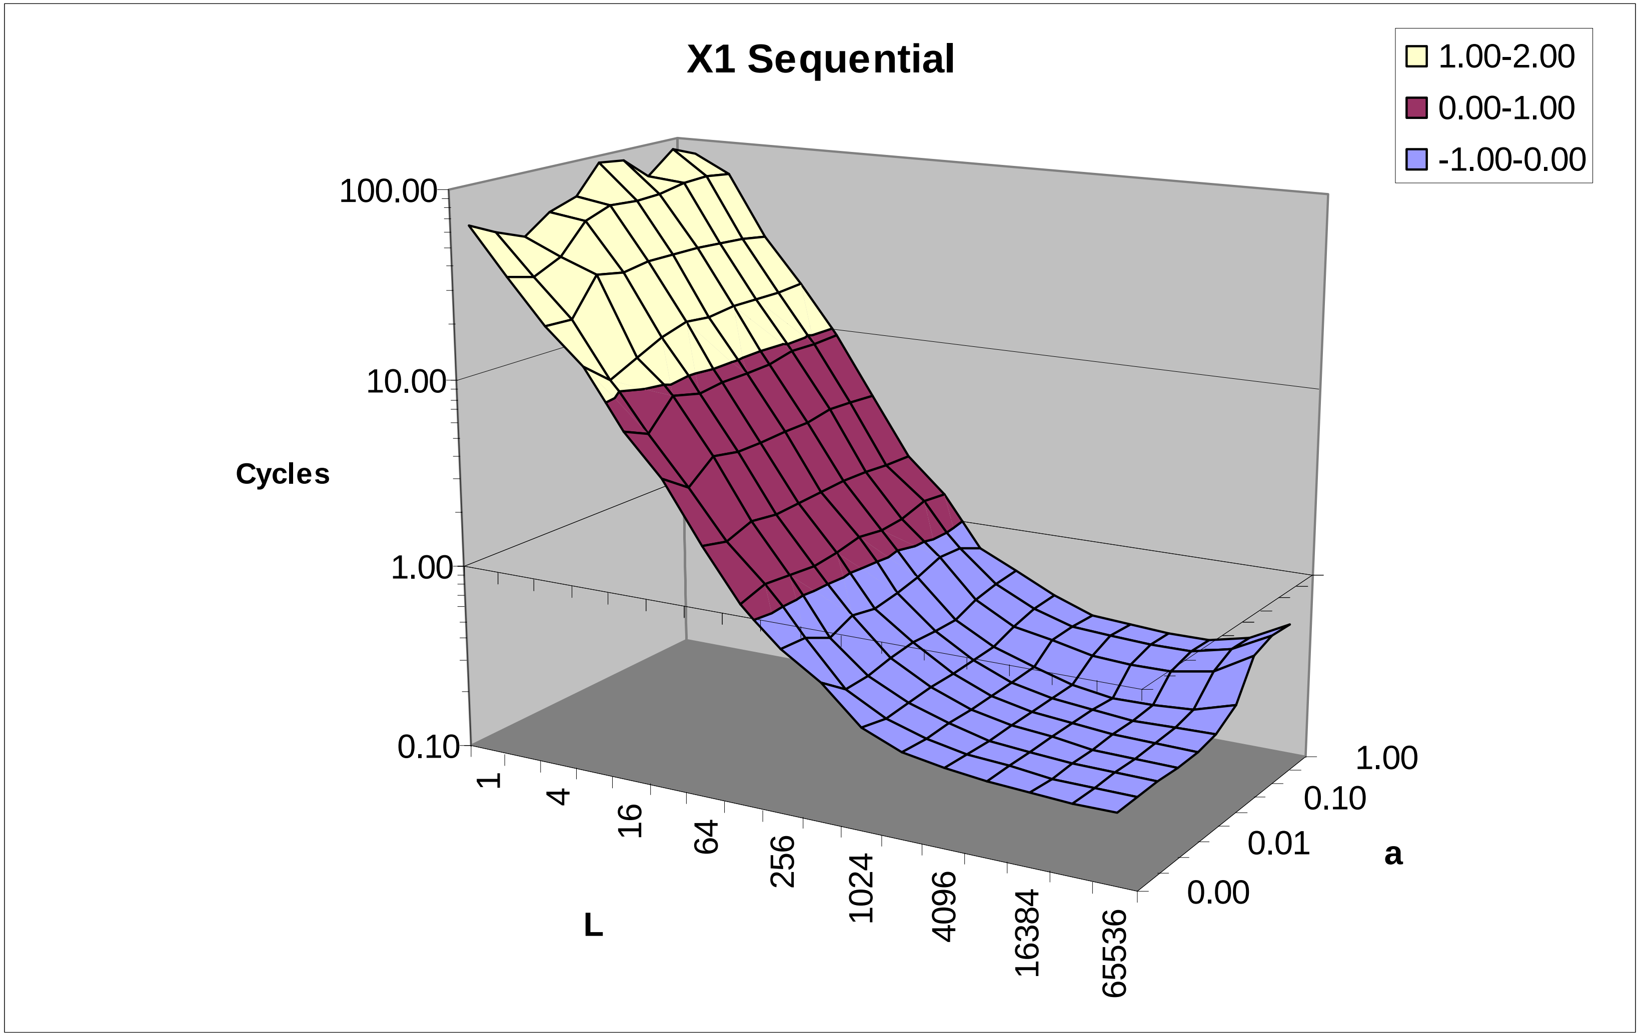

Apex-Map Sequential

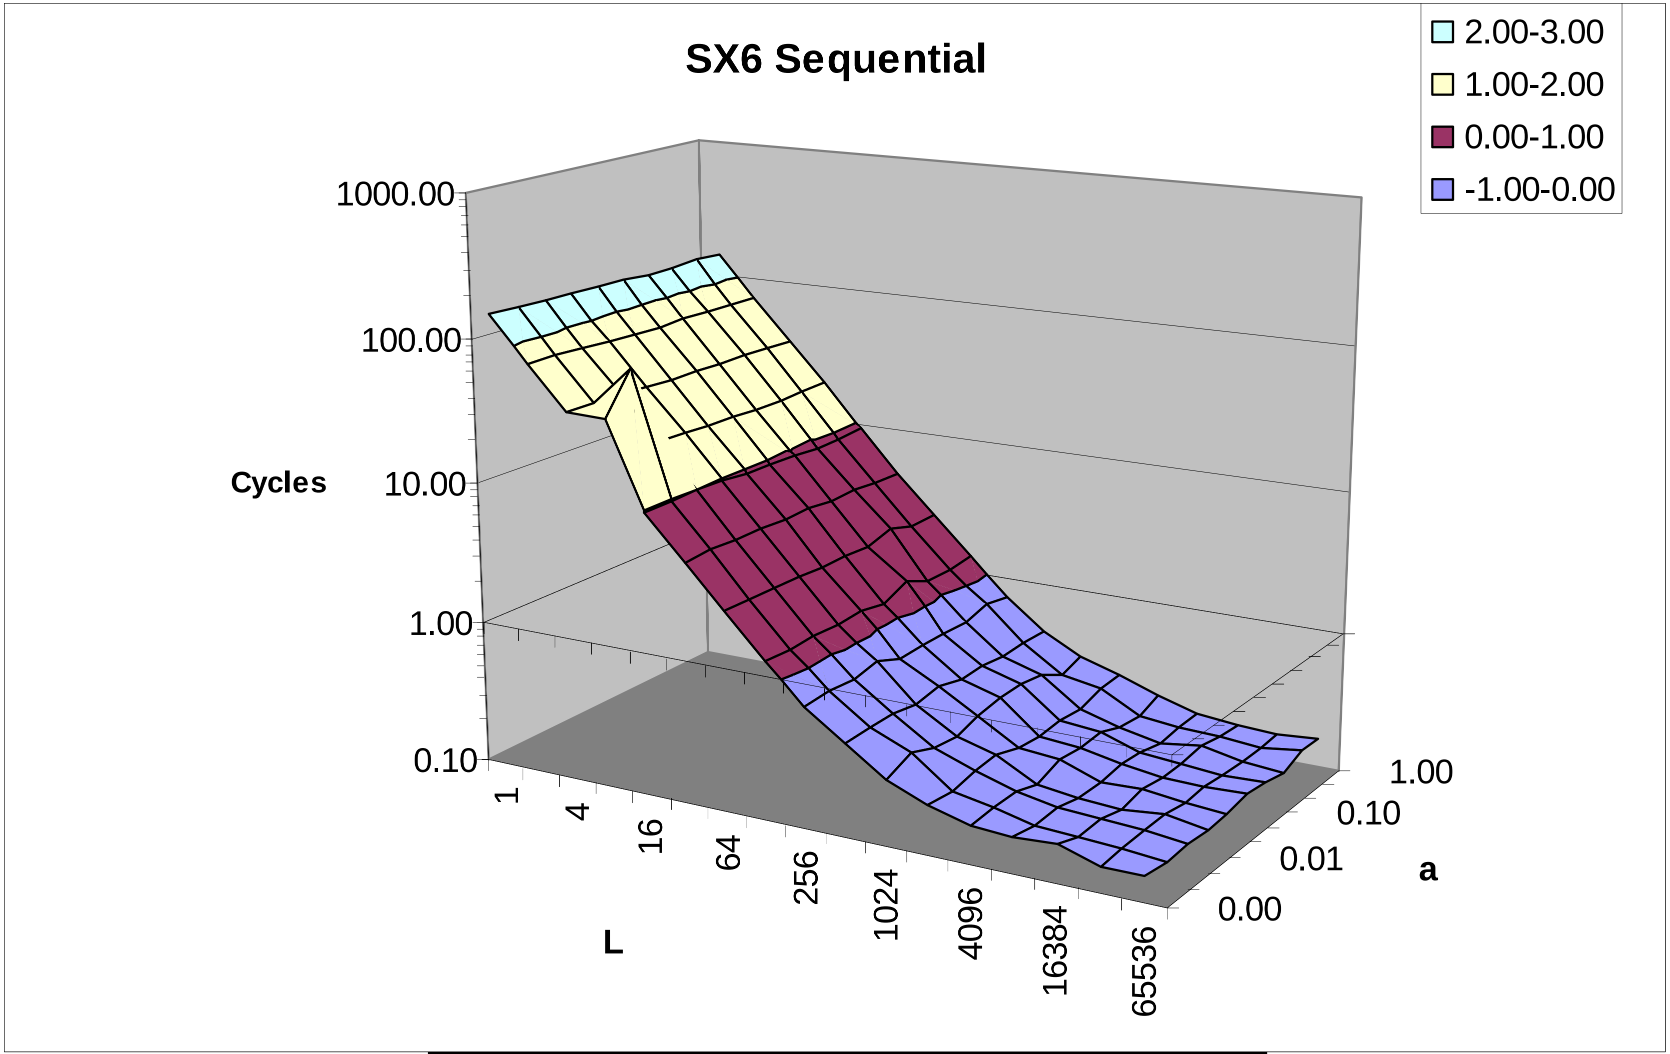

Parallel APEX-Map

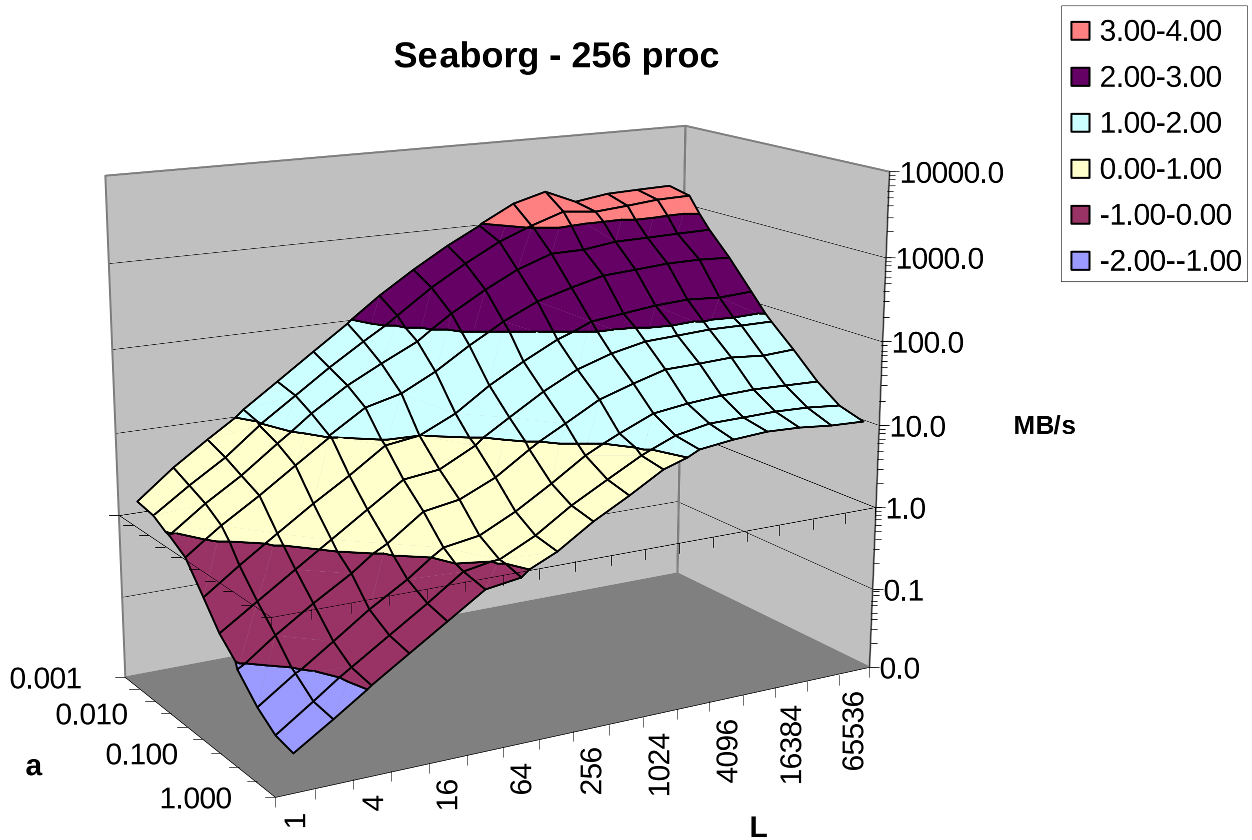

Parallel APEX-Map

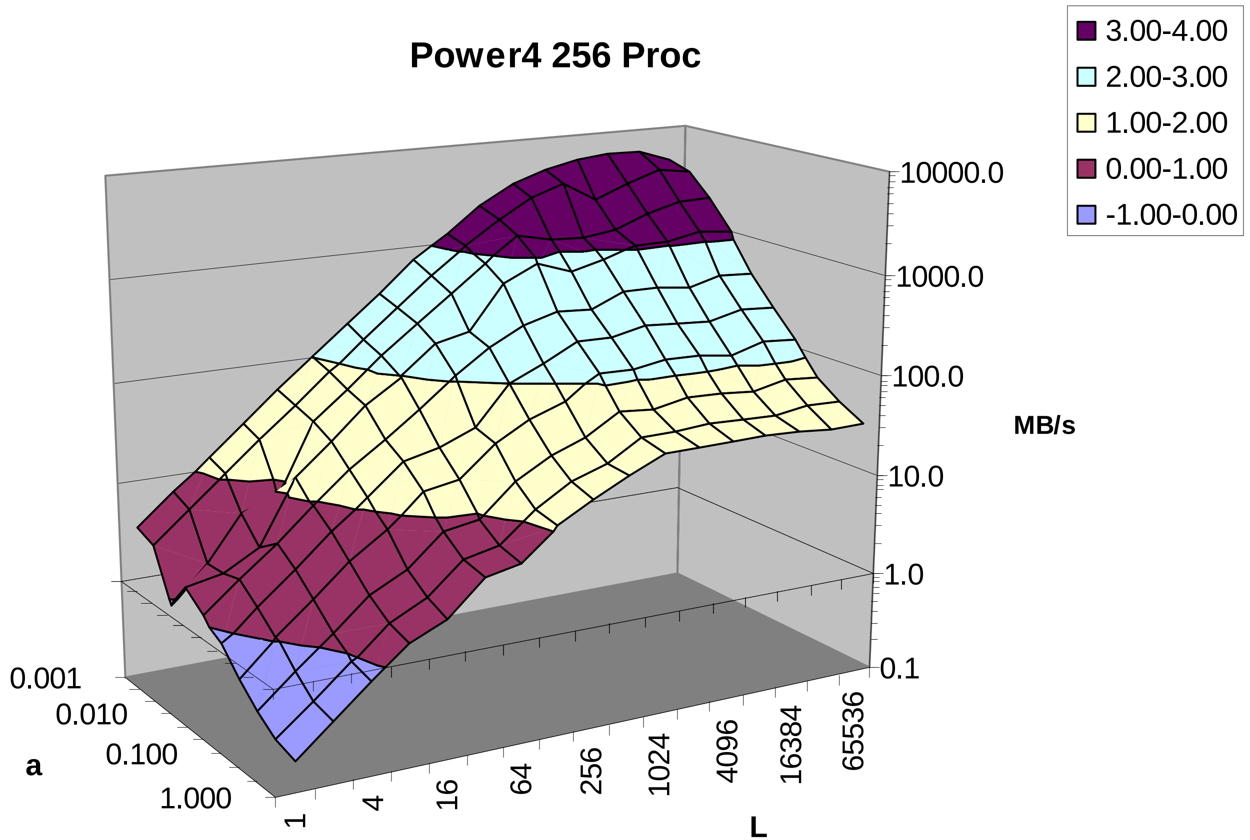

Parallel APEX-Map

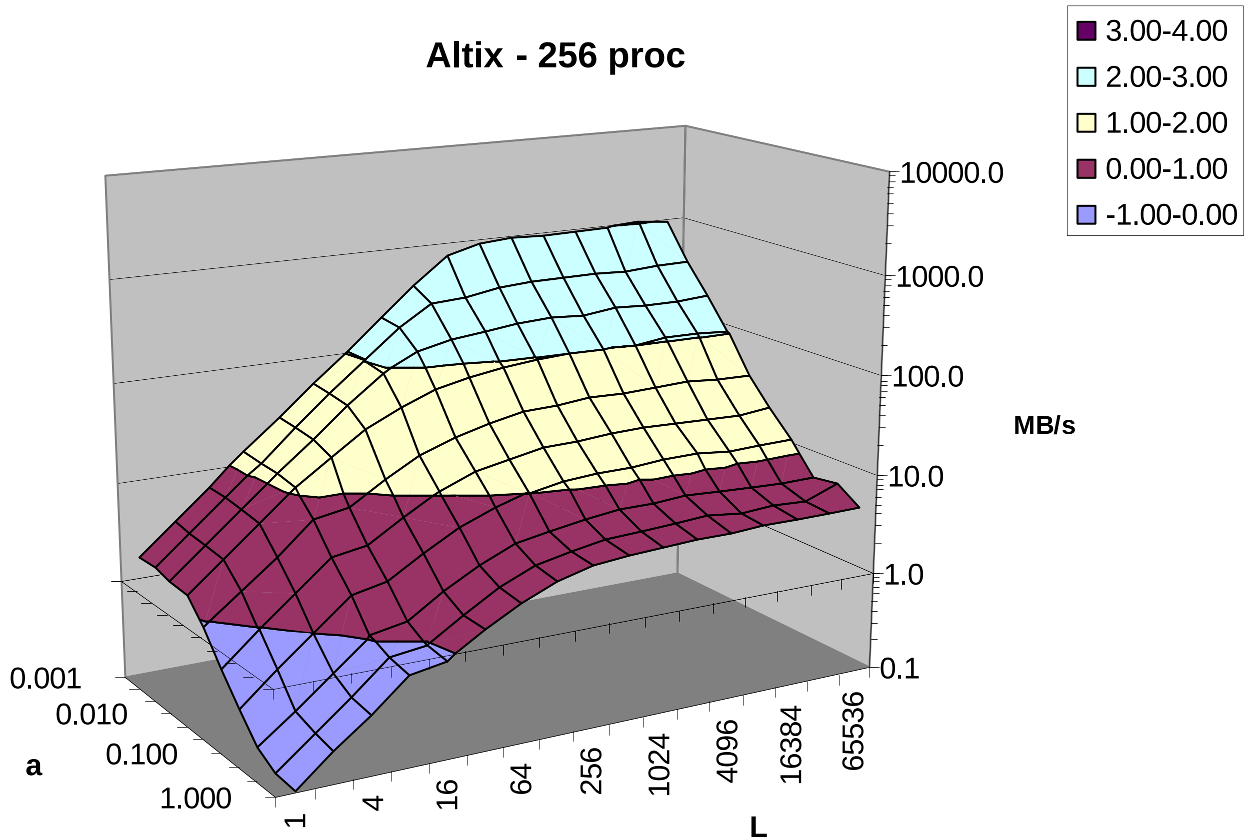

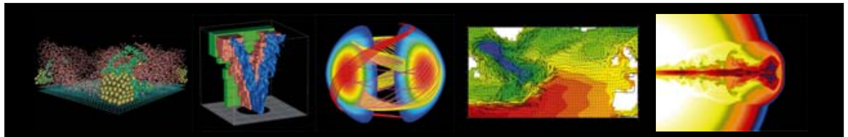

Parallel APEX-Map

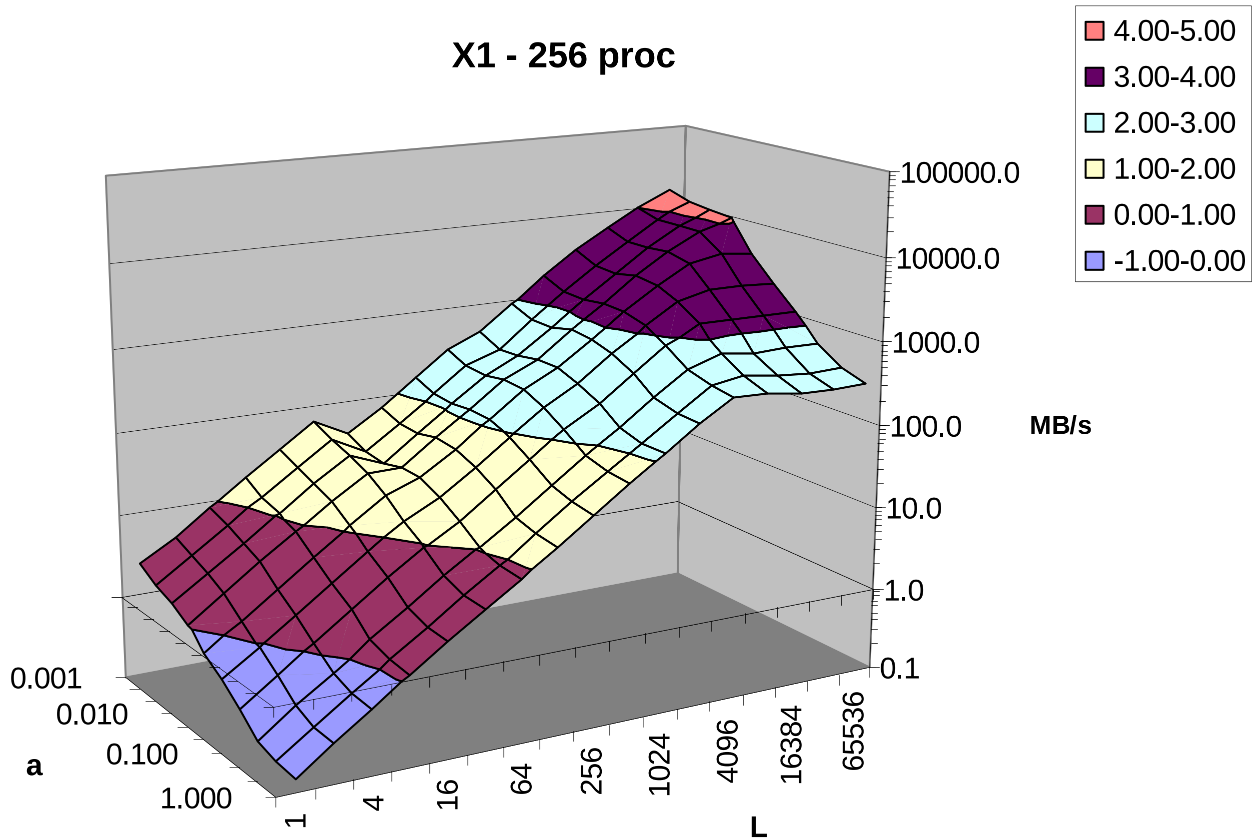

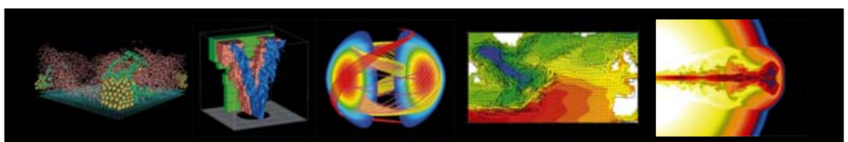

Parallel APEX-Map

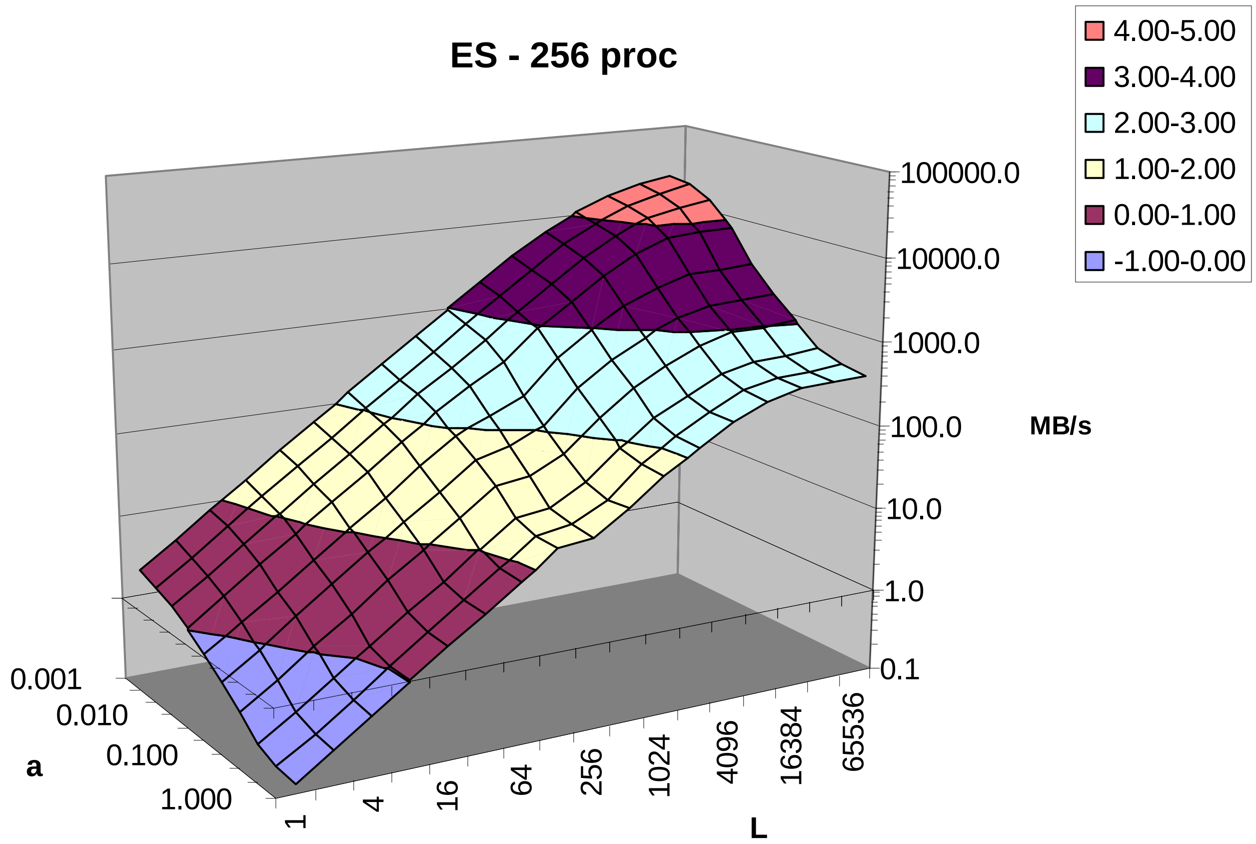

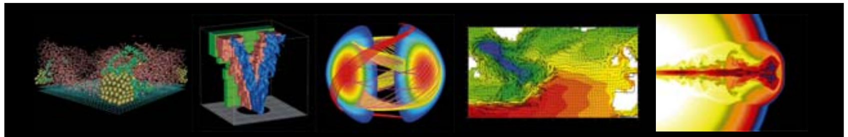

BERKELEY LAB

Office of Science
U.S. DEPARTMENT OF ENERGY

Summary

- Progress in applications performance on leadership computing platforms requires development of sciences driven architecture
- In 2004 we see the emergence of a quantitative, scientific approach to applications characterization on advanced platforms